



Standards HEC-IWG

Tom Ruwart
August 16, 2005



Overview

- 50,000 foot view of BIG animals
- Relevant Standards Organizations and Efforts
 - SNIA
 - ANSI T10, T11, T13
 - IETF
 - RDMA Consortium
 - OpenIB
 - Myrinet
- What this has to do with Research



SNIA

- Object-based Storage Device (OSD)
Technical Working Group (TWG)
 - Revision 1 in use
 - Revision 2 in development
- Storage Management – SMI
- Other potential standard definitions



ANSI T10

- SCSI Protocol definitions
- SAS – Serial Attached SCSI
 - Released and in production
 - Adoption seems to be gaining significant momentum for array infrastructures
- OSD Protocol
 - Rev 1 released and early deployments begun
 - Currently working on Rev 2 (SNIA)



ANSI T11

- Everything Fibre Channel
 - 2Gig FC currently shipping
 - 4Gig FC in early deployment
 - 10Gig FC on the horizon
- Still the storage connection of choice in the enterprise space



ANSI T13

■ Everything ATA

□ Parallel ATA

- Current standard for consumer-grade disk drives
- Being replaced by Serial ATA on most 3.5-inch disks and some 2.5-inch disks
- Still used for very small form-factor drives

□ Serial ATA

- Gaining significant momentum and widely available
- Used for high-density packaging of consumer-grade disk drives
- Includes many of the features previously afforded only to FC
- Currently at 1.5 Gb/s, moving to 3Gb/s



IETF

- Ethernet – 10Gig
- NFSv4 (pNFS)
- IP Storage
 - iSCSI – in wide deployment
 - iFCP & FCIP – deployed
 - iSER, iSNS, DA – in development



RDMA Consortium

- Remote DMA over Ethernet or any transport
- Sockets Direct Protocol
- iSCSI extensions for RDMA (iSER)
 - Version 1.0 released
- Data-mover Architecture (DA)
 - Version 1.0 released



InfiniBand

- Still a niche market in cluster interconnects
- Gaining some traction as a storage interconnect
- OpenIB - a source for IB software for Linux
- Future is unclear – possibly eclipsed by competing technologies such as Myrinet, Quadrics, PCIExpress/AS



Myrinet

- Myrinet is the most widely adopted cluster interconnect
- Myrinet transport is an ANSI standard
- Myricom provides proprietary protocols for value-add



What this has to do with Research

- These technologies are continually evolving
- Need insight about
 - How fast they evolve
 - The direction of evolution
 - Potential barriers
 - Other influences on evolution – markets, ecosystems, new technologies
- These are independent foundation technologies
- Interdisciplinary research
 - Clustered disk drives
 - Intelligent, autonomous storage systems
 - Transparently-hybrid storage systems